

# Mind2Drive: Predicting Driver Intentions from EEG in Real-world On-Road Driving

Ghadah Alosaimi<sup>1,2</sup>, Hanadi Alhamdan<sup>4,2</sup>, Wenke E<sup>2</sup>, Stamos Katsigiannis<sup>2</sup>,  
Amir Atapour-Abarghouei<sup>2</sup>, Toby P. Breckon<sup>2,3</sup>

<sup>1</sup>Department of Computer Science, Imam Mohammad Ibn Saud Islamic University, Saudi Arabia  
Department of {<sup>2</sup>Computer Science, <sup>3</sup>Engineering}, Durham University, UK

<sup>4</sup>Department of Computer Science, Princess Nourah bint Abdulrahman University, Saudi Arabia

**Abstract**—Predicting driver intention from neurophysiological signals offers a promising pathway for enhancing proactive safety in advanced driver assistance systems, yet remains challenging in real-world driving due to EEG signal non-stationarity and the complexity of cognitive–motor preparation. This study proposes and evaluates an EEG-based driver intention prediction framework using a synchronised multi-sensor platform integrated into a real electric vehicle. A real-world on-road dataset was collected across 32 driving sessions, and twelve deep learning architectures were evaluated under consistent experimental conditions. Among the evaluated architectures, TSception achieved the highest average accuracy (0.907) and Macro-F1 score (0.901). The proposed framework demonstrates strong temporal stability, maintaining robust decoding performance up to 1000 ms before manoeuvre execution with minimal degradation. Furthermore, additional analyses reveal that minimal EEG preprocessing outperforms artefact-handling pipelines, and prediction performance peaks within a 400–600 ms interval, corresponding to a critical neural preparatory phase preceding driving manoeuvres. Overall, these findings support the feasibility of early and stable EEG-based driver intention decoding under real-world on-road conditions. Code: <https://github.com/galosaimi/Mind2Drive>.

**Index Terms**—Brain-machine Interfaces, Neural Network Applications, Industrial applications, Transformer Networks.

## I. INTRODUCTION

Accurate driver intention prediction is essential for road safety and effective interaction between human-driven vehicles and driver assistance or autonomous systems. Critical manoeuvres such as turning or braking are preceded by internal decision-making processes that are not directly observable from vehicle motion. Recent studies identify limitations in understanding human driving intentions as a major barrier to safe human–machine coexistence [1]. Even modest gains in anticipation time, on the order of hundreds of milliseconds, can benefit collision avoidance and cooperative driving [2].

Most driver intention prediction methods rely on external observations such as vehicle kinematics, trajectories, and traffic interactions. Machine learning and physics-informed models using speed, yaw rate, lane position, and inter-vehicle distances achieve high accuracy for manoeuvres such as lane changes and lane keeping [3]. However, because these methods depend on observable post-manoevr motion cues, performance degrades with longer prediction horizons [3], [4].

Electroencephalography (EEG) offers a distinct sensing modality by directly capturing neural activity related to percep-

tion, decision-making, and motor preparation. Prior research has shown that intention-related neural signatures precede physical movement, which enables early motor intention decoding from EEG signals [5]. Recent advances in deep learning further demonstrate that complex and fine-grained motor intentions can be decoded from EEG in real time, even under noisy conditions [6].

In the driving domain, EEG-based intention prediction has been studied for tasks such as emergency braking, steering, and acceleration, demonstrating that EEG can anticipate driving intentions earlier than vehicle-based signals [2], [7]. However, most studies are conducted in simulators or controlled settings, and real-world on-road EEG experiments remain limited, with inconsistent preprocessing, labelling, and evaluation [8].

Another critical challenge lies in EEG preprocessing. EEG signals have inherently low signal-to-noise ratio and are strongly affected by motion artefacts and environmental interference in real-world driving. Comparative studies show that preprocessing choices substantially influence decoding performance, with no single pipeline being universally optimal [9]. Nevertheless, preprocessing is often treated as a fixed design choice rather than as a systematic experimental factor.

Motivated by these gaps, this study investigates deep learning models, including transformer-based architectures, for EEG-based prediction of driver intentions during real-world on-road driving. Using synchronised EEG and high-precision vehicle kinematics, intentions are defined via motion-based labelling and predicted across multiple future horizons. In contrast to prior work, this study systematically evaluates EEG preprocessing pipelines and quantifies how early intentions can be decoded under realistic conditions, thereby advancing neuro-informed driver intention prediction for practical intelligent transportation systems.

## II. RELATED WORK

We consider related work across the areas of vehicle-dynamics–based intention prediction (Section II-A), EEG motor-intention decoding (Section II-B), EEG-based driving intention prediction (Section II-C), human-centred driving systems using physiological cues (Section II-D), and methodological issues specific to real-world EEG (Section II-E).

### A. Driver Intention Prediction Using Vehicle Dynamics

Prior studies infer driver intentions from vehicle-centric observations by modelling temporal patterns in kinematic and interaction variables. Time-sequenced and trajectory-based approaches, such as weighted hidden Markov models, have shown effective recognition of lane-change and lane-keeping behaviours from motion histories [3]. More recent physics-informed and interaction-aware learning frameworks further improve robustness and accuracy by incorporating vehicle dynamics and traffic constraints [4]. However, because these methods infer intention from observable motion, their predictions remain closely coupled to kinematic onset, limiting characterisation of earlier, pre-movement intention stages.

### B. EEG-Based Intention and Motor Preparation

Idowu et al. formulated motor intention recognition as a multi-class classification problem using fixed-length EEG segments aligned to explicitly cued upper-limb movements, achieving improved accuracy through deep architectures that combine spatial-spectral feature extraction with temporal modelling [5]. This evaluation relies on segmented trials with known intention onset, assuming externally defined task timing. More recent work demonstrated real-time decoding of individual finger movements using end-to-end deep learning networks trained on densely supervised motor tasks, achieving low latency and high classification stability under laboratory conditions [6]. However, these approaches operate in tightly designed laboratory tasks in which road events, action timing, and movement boundaries are predefined, and subjects perform clearly separated, labelled movements in response to explicit cues (e.g., pressing a button, moving a finger, or lifting an arm). Consequently, they do not address the continuous, self-initiated intention formation or multi-horizon prediction required in real-world driving.

### C. EEG-Based Driver Intention Prediction

Several studies have explored EEG-based intention prediction in driving, mainly focusing on braking, steering, or acceleration under controlled conditions. Liang et al. formulated emergency braking detection as a binary classification problem using EEG segments aligned to braking onset, demonstrating anticipation of braking actions in simulated driving [2]. Li et al. combined EEG and galvanic skin response signals to predict multiple driving intentions via multimodal fusion, while still relying on predefined manoeuvre events and simulator-based evaluation [7]. More recent work applied few-shot transfer learning on neuromorphic hardware for individualised braking intention detection, but was limited to braking-only and structured tasks [10]. Overall, these approaches focus on discrete, event-driven intentions and do not address continuous or multi-horizon prediction in real-world on-road driving.

### D. Human-Centred Driving Systems

Recent work has incorporated human behavioural and neural signals as auxiliary inputs for autonomous driving, rather than directly decoding driver intention. Xia et al. integrated

synchronised human behaviour data, extracted from EEG and other physiological signals, into end-to-end driving models in simulated environments, demonstrating that such signals can influence policy learning under controlled conditions [1]. However, EEG was used as indirect supervision and evaluation was limited to simulation, without addressing explicit intention decoding or early prediction in real-world driving [11].

### E. Methodological Challenges and Real-World EEG

A key challenge in EEG-based intention prediction studies is the robustness of signal preprocessing under real-world conditions. Coelli et al. showed that preprocessing choices can significantly affect decoding performance, with no single pipeline universally optimal across tasks and datasets [9]. Tao et al. further highlighted that the limited availability of realistic large-scale driving datasets incorporating physiological signals constrains evaluation under realistic conditions [8].

Vehicle-centric predictors remain tethered to observable manoeuvres, while EEG-based methods largely stay confined to simulator/episodic paradigms, which leaves continuous intention decoding in real on-road settings underexplored.

## III. METHODOLOGY

This section describes the experimental platform, data collection protocol, problem formulation, and deep learning models used in this study.

### A. Sensor Platform and Data Collection

Data were collected using a Renault Twizy electric vehicle instrumented with a synchronised multi-sensor platform for real-world driving experiments. Vehicle layout and sensor placement are shown in Fig. 1, and sensor specifications are summarised in TABLE I. All sensor streams were recorded using the Robot Operating System (ROS Noetic), providing unified timestamping for precise temporal alignment across data sources.

Vehicle motion was measured using a GNSS/INS unit with RTK correction via NTRIP, yielding accurate estimates of vehicle position, planar velocity, yaw angle, and yaw rate, which were used to derive ground-truth driving actions and intention labels. Additional sensors, including a stereo camera and LiDAR with an integrated IMU, were installed to provide visual and inertial context and support synchronisation. It should be noted that data from these sensors were not used as model inputs, but have been reserved for future work.

EEG signals were recorded using a wireless 16-channel OpenBCI Cyton and Daisy system with a gel-free electrode cap configured according to the international 10–20 system [12]. The cap provided coverage over frontal, central, temporal, parietal, and occipital regions, with electrode locations listed in TABLE I.

Four participants (three male and one female, aged 25–45 years), each holding a valid UK driving licence, participated in the study. All had normal or corrected-to-normal vision and reported no history of neurological or psychiatric disorders.

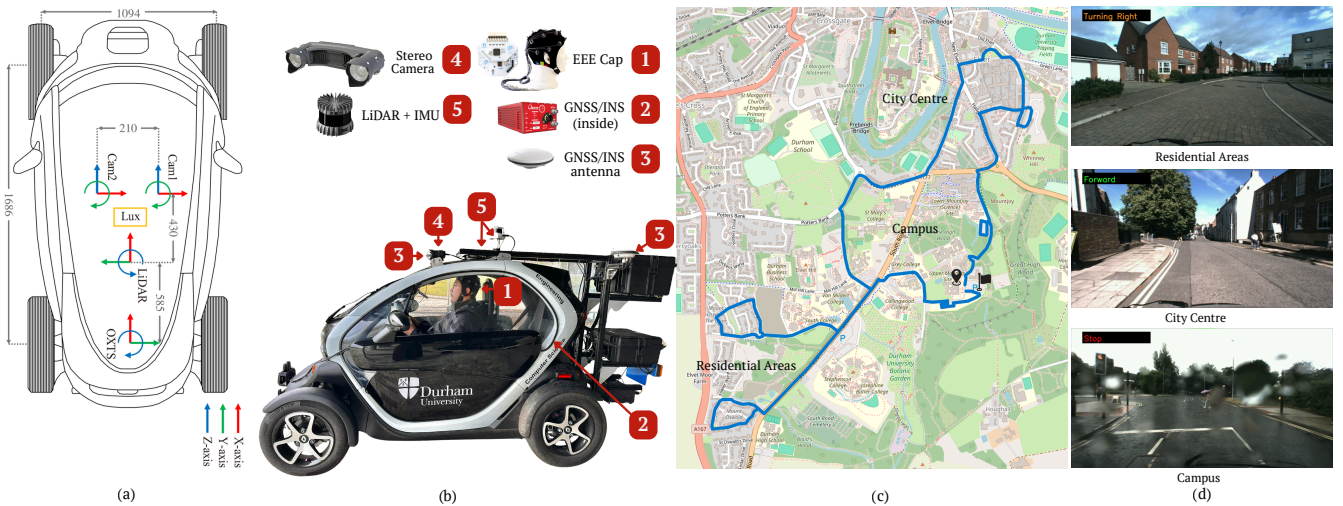


Fig. 1: Experimental platform and data collection environment. (a, b) Instrumented Renault Twizy showing sensor placement and coordinate frames, including EEG, GNSS/INS, stereo cameras, and a LiDAR–IMU, with all sensor streams synchronised via ROS and recorded in rosbag format. (c) Predefined driving route across campus, city-centre, and residential areas. (d) Forward-facing camera views illustrating data collection under different environmental conditions.

TABLE I: Sensor specifications and data characteristics

Sensor Type	Model	Sampling Rate	Channels / Beams	Measured Parameters
EEG	Gel-free OpenBCI cap (Cyton + Daisy)	125 Hz	16	Scalp EEG (Fp1, Fp2, C3, C4, T3, T4, O1, O2, F7, F8, F3, F4, T5, T6, P3, P4)
GPS / INS	OxTS RT3000v3	100 Hz	2 antennas	Position (latitude, longitude), velocity ( $v_x, v_y$ ), yaw ( $\psi$ ), yaw-rate ( $\dot{\psi}$ )
LiDAR + IMU	Ouster OS1-128	10 Hz	128	3D point clouds, angular velocity, linear acceleration
Cameras	Carnegie Robotics Multisense S21 (stereo)	30 Hz	2 (left/right)	RGB images (2048 × 1088 pixels)

Written informed consent was obtained from all participants in accordance with institutional ethical guidelines.

Data collection took place on public roads in Durham, UK, covering campus, city-centre, and residential environments under varying weather conditions, including sunny, cloudy, and rainy. All participants followed the same predefined route of approximately 9 km to ensure comparable driving conditions while preserving realistic variability in traffic and driving behaviour. Representative environments and the driving route are shown in Fig. 1.

Each participant completed 8 driving sessions, yielding a total of 32 sessions. Each session involved a single traversal of the route with an average duration of approximately 22 minutes. Prior to data collection, participants completed two familiarisation sessions to acclimate to the vehicle, route, and experimental setup. Participants were instructed to drive naturally while complying with traffic regulations. All sessions meeting data-quality criteria, stable EEG recordings with electrode impedance below 10 k $\Omega$  and complete kinematic data without extended dropouts, were retained for analysis.

### B. Signal Synchronisation and Preprocessing

EEG signals sampled at 125 Hz were used as the reference timeline for synchronisation within the ROS framework. Vehicle kinematic signals from the GNSS/INS system were

linearly interpolated to the EEG timestamps, ensuring temporal alignment between neural activity and vehicle motion.


Prior to model input, EEG signals were normalised to ensure consistency across sessions and participants. No automated preprocessing pipeline was applied by default. To assess the impact of preprocessing, commonly used EEG cleaning strategies were subsequently evaluated, including PyPREP-based pipelines with band-pass filtering and optional RANSAC-based bad channel detection [13]. Additional artifact handling approaches combining independent component analysis (ICA), artifact subspace reconstruction (ASR), and automated trial rejection were also tested [14], [15]. All configurations were applied consistently across models and sessions, and their effects are analyzed in Section IV-B.

### C. Motion-Based Action Labelling


Ground-truth driving actions were derived from vehicle kinematics measured by the GNSS/INS system. Continuous motion signals were converted into discrete manoeuvre labels based on vehicle speed and yaw-rate. Vehicle speed at time  $t$  was computed from the planar velocity components as:

$$v(t) = \sqrt{v_x^2(t) + v_y^2(t)}, \quad (1)$$


where  $v_x(t)$  and  $v_y(t)$  denote the forward–backward and lateral velocity components in the road plane. Vehicle orientation




(a) Vehicle speed over time with motion labels.



(b) Yaw-rate over time with turning thresholds.



(c) Kernel density estimate of vehicle speed.



(d) Kernel density estimate of yaw-rate magnitude.

Fig. 2: Motion-based action labelling from vehicle kinematics. Representative time-series of speed and yaw rate with applied thresholds for a single session of a subject (top), and kernel density estimates across all sessions validating the selected stop and turning thresholds (bottom).

was described by the yaw angle  $\psi(t)$ , and turning dynamics by the yaw rate  $\dot{\psi}(t)$ .

To distinguish forward and reverse motion, the direction of movement was computed as:

$$\theta_v(t) = \text{atan2}(v_y(t), v_x(t)), \quad (2)$$

and compared with the vehicle heading  $\psi(t)$ . The angular difference

$$\Delta\theta(t) = \text{wrap}(\theta_v(t) - \psi(t)) \quad (3)$$

was used to identify forward ( $|\Delta\theta(t)| \leq \pi/2$ ) and reverse motion ( $|\Delta\theta(t)| > \pi/2$ ).

Two thresholds were defined: a speed threshold  $v_{th}$  to separate stationary and moving states, and a yaw-rate threshold  $\omega_{th}$  to distinguish straight driving from turning. Vehicle motion at time  $t$  was classified as:

- Stopped:**  $v(t) < v_{th}$
- Turning Left:**  $v(t) \geq v_{th}$  AND  $\dot{\psi}(t) > \omega_{th}$
- Turning Right:**  $v(t) \geq v_{th}$  AND  $\dot{\psi}(t) < -\omega_{th}$
- Moving Straight:**  $v(t) \geq v_{th}$  AND  $|\dot{\psi}(t)| \leq \omega_{th}$  AND  $|\Delta\theta(t)| \leq \frac{\pi}{2}$
- Reverse:**  $v(t) \geq v_{th}$  AND  $|\Delta\theta(t)| > \frac{\pi}{2}$

Thresholds were selected through visual inspection of representative time-series and validated using global distributions across all sessions. As shown in Fig. 2, speed and yaw-rate signals exhibit clear separation between stationary, straight-driving, and turning regimes, which is further confirmed by kernel density estimates. The resulting vehicle motion labels were additionally validated through manual visual inspection of the recorded videos for a subset of the samples. Reverse and stopped states were excluded from the final action set, as reverse manoeuvres were not part of the predefined route and the study focuses on continuous forward-driving behaviour.

#### D. Intention Prediction Problem Formulation

This study aims to predict a driver's future driving intention from EEG signals prior to manoeuvre execution. The problem is formulated as a supervised task in which neural activity preceding a manoeuvre is used to predict a future action.

Let  $X(t) \in \mathbb{R}^{C \times T}$  denote the multichannel EEG signal observed up to time  $t$ , where  $C$  is the number of EEG channels and  $T$  is the number of temporal samples in the observation window. Let  $y(t)$  denote the discrete driving action being executed by the vehicle at time  $t$ , obtained from motion-based labelling. A future-intention label at prediction horizon  $\Delta$  is defined as

$$y_{\Delta}(t) = y(t + \Delta), \quad (4)$$

with  $\Delta \in \{100, 200, 300, \dots, 1000\}$  ms. The learning objective is to map an EEG window  $X(t - w : t)$  to the corresponding future action  $y_{\Delta}(t)$ , where  $w$  denotes the window length. The prediction task is restricted to the action set:  $y \in \{\text{Forward, Turn-Left, Turn-Right}\}$ .

Multiple prediction horizons are considered to quantify how early driver intentions can be decoded from EEG. Short horizons (100–300 ms) correspond to near-term motor preparation, while longer horizons (500–1000 ms) reflect earlier stages of cognitive intention formation.

#### E. Experimental Design and Evaluation Protocol

Twelve deep learning architectures were evaluated to provide a comprehensive assessment of EEG-based driver intention classification across diverse model families. The selected models span convolutional, recurrent, and attention-based paradigms. Convolutional models extract local spatiotemporal EEG patterns through temporal and spatial filtering [16], and include TSception [17], ShallowConvNet and DeepConvNet [18], EEGNet [19], CCNN [20], DGCNN [21], CNN1D [22], and STNet [23]. Recurrent models capture

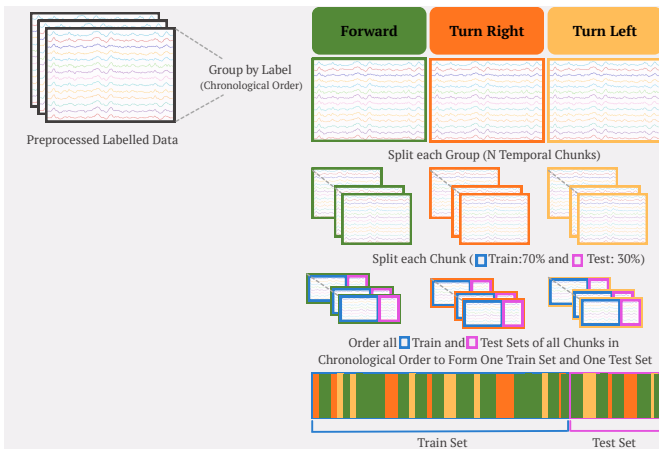


Fig. 3: EEG recording is grouped by class and ordered chronologically. Each group is split into 70%-train and 30%-test segments using temporal chunks, then re-ordered to form the final train and test sets.

sequential EEG dependencies [24] and are represented by LSTM and GRU [25]. Attention-based architectures combine convolutional feature extraction with mechanisms for modelling long-range temporal relationships [26], and include EEGConformer [27] and ViT [28]. This selection was designed to cover a broad spectrum of architectural inductive biases, temporal modelling capabilities, and spatial feature extraction strategies commonly adopted in EEG decoding research.

Furthermore, based on preliminary experimentation and the prohibitively large number of possible configuration combinations, four representative models, EEGNet, ShallowConvNet, GRU, and EEGConformer, were selected for an ablation study, as they reflect complementary architectural families. The ablation study focused on identifying optimal data preparation strategies, EEG preprocessing pipelines, and the impact of using 8 versus 16 EEG channels.

All models were trained and evaluated using identical data splits and windowing configurations to ensure that performance differences reflect preprocessing or architectural effects rather than data preparation variability. To ensure fair comparison, standardised training settings were used across models. A batch size of 128 and a maximum of 2,000 training epochs were applied, with the Adam optimiser; all models used a learning rate of 0.001 except EEGConformer ( $1e-4$ ), with all learning rates remaining in the same order of magnitude as those reported in their original papers. Loss functions were selected according to model output formulations, with class-weighted variants used where necessary to address class imbalance. The model checkpoint achieving the best Macro-F1 score during held-out validation was selected for final evaluation. Data splitting and training were performed at the session level to preserve temporal continuity and avoid information leakage. All experiments were conducted on multiple NVIDIA RTX A6000 GPUs to ensure consistent computational conditions.

Performance was primarily assessed using Macro-F1 to account for class imbalance between forward-driving and turning manoeuvres. Balanced accuracy and per-class recall were also reported to provide complementary performance insights.

## IV. EXPERIMENTAL RESULTS

This section presents a systematic evaluation of the proposed EEG-based driver intention prediction framework, assessing the effects of data preparation strategies, preprocessing pipelines, and model architectures.

### A. Evaluation of data preparation strategies

The impact of data preparation choices on EEG-based driver intention prediction performance is analyzed prior to preprocessing and model-specific analysis, with the goal of identifying a robust strategy for subsequent experiments.

A systematic ablation was conducted over data splitting, windowing, and label aggregation strategies. To prevent temporal leakage, data were split on continuous and time-ordered EEG signals before windowing. Windowing was applied independently within training and test sets, and oversampling, when used, was restricted to the training data. Twelve windowing configurations were evaluated by combining four window lengths (0.5s, 1.0s, 2.0s, and 5.0s) with two step sizes (non-overlapping and 50% overlap). Window-level labels were assigned using either a majority rule, which assigns the most frequent action within a window, or a rejection rule, which retains only windows containing a single, consistent manoeuvre and discards windows spanning action transitions, thereby ensuring manoeuvre purity. The rejection rule was applied only to 0.5s and 1.0s windows, as longer windows showed excessive rejection rates. Each configuration was tested using two data-splitting strategies: a standard temporal split (70%/30%) and a label-stratified split preserving class balance while maintaining temporal order, as illustrated in (Fig. 3). Three imbalance-handling settings were also considered: no oversampling, random oversampling, and TimeGAN-based augmentation [29].

Performance was evaluated using average Macro-F1 and balanced accuracy across the four models selected for the ablations to reduce architecture-specific bias, with only top configurations reported. Across all settings, a 1.0s window with 50% overlap and the rejection labelling rule achieved the best performance, outperforming majority voting. The choice of data split had a substantial impact on performance. The label-stratified split significantly outperformed the standard split, achieving an average accuracy of 73.88% compared to 36.53%, corresponding to an improvement of approximately 37%. This highlights the importance of preserving class balance between training and test sets for reliable intention prediction. On average, 27.1% of windows per session were rejected due to ambiguous labels, reducing the mean number of windows from approximately 5,096 to 3,731 per session. Despite this reduction, rejecting ambiguous windows improved label consistency and downstream classification performance, with the rejection rule outperforming majority voting by approximately 11%.

The effect of class imbalance handling is summarised in TABLE II. Random oversampling achieved the highest average Macro-F1 (0.873) across four models, outperforming both the baseline configuration without oversampling (0.828) and TimeGAN-based augmentation (0.811), indicating improved

TABLE II: Comparison of the top-performing data preparation strategies, selected from all evaluated combinations, reported using the average metrics across the four models used for ablations.

Win.	Ov.	Agg.	Sampling Strategy	Macro-F1	Bal. Acc
<b>1.0s</b>	<b>50%</b>	<b>Reject</b>	<b>Random</b>	<b><math>0.873 \pm 0.038</math></b>	$0.838 \pm 0.044$
1.0s	0%	Reject	Random	$0.855 \pm 0.042$	$0.812 \pm 0.044$
1.0s	50%	Reject	None	$0.828 \pm 0.059$	$0.858 \pm 0.037$
1.0s	50%	Reject	TimeGAN	$0.811 \pm 0.048$	$0.852 \pm 0.027$
0.5s	25%	Reject	Random	$0.806 \pm 0.049$	$0.780 \pm 0.056$

Note: Win. = window length; Ov. = overlap; Agg. = aggregation; Bal. Acc = Balanced Accuracy; None indicates no sampling (baseline).

performance on minority driving intention classes. While the baseline configuration achieved the highest average balanced accuracy (0.858), its Macro-F1 was lower than that of random oversampling. TimeGAN achieved competitive balanced accuracy (0.852), but did not consistently outperform random oversampling, suggesting limited benefit for this task.

Based on this ablation, a label-stratified split with random oversampling, combined with 1.0s windows, 50% overlap, and the rejection labelling rule, was selected as the fixed data preparation strategy for all subsequent experiments.

### B. Evaluation of EEG Preprocessing Pipelines

The effect of commonly used EEG preprocessing pipelines, described in Section III-B, on model performance is evaluated using the above fixed data preparation strategy, with the aim of assessing whether advanced preprocessing offers benefits over minimal preprocessing. Results in TABLE III show a consistent trend across all models: minimal preprocessing achieved the best overall performance. The baseline configuration using only z-score normalisation (without PyPREP) yielded the highest average accuracy (0.839) and Macro-F1 (0.873), indicating that raw EEG signals already contain discriminative information for driver intention prediction. Applying PyPREP without RANSAC produced comparable results but did not improve performance. In contrast, pipelines incorporating aggressive artefact handling, such as RANSAC-based channel rejection, ICA, ASR, and automated trial rejection consistently degraded performance, suggesting that intention-related neural information may be removed alongside noise in real-world EEG.

Based on these results, the baseline preprocessing configuration without PyPREP was selected for all subsequent experiments to maximise performance while maintaining simplicity and reproducibility.

### C. Effect of Spatial EEG Resolution (8 vs. 16 Channels)

To evaluate the impact of spatial resolution, models were trained using both 8-channel and 16-channel EEG configurations under identical data preparation and preprocessing settings (TABLE IV). For the 8-channel configuration, electrodes F3, F4, F7, F8, C3, C4, P3, and P4 were selected based on established neurophysiological evidence linking frontal regions to executive control and response inhibition, central regions to motor preparation, and parietal regions to spatial attention and

TABLE III: Comparison of different EEG preprocessing pipelines, reported using the average metrics across the four models used for ablations.

Preprocessing Pipeline	Macro-F1	Bal. Acc
<b>No PyPREP (Baseline)</b>	<b><math>0.873 \pm 0.038</math></b>	<b><math>0.839 \pm 0.047</math></b>
PyPREP (No RANSAC)	$0.823 \pm 0.062$	$0.831 \pm 0.060$
AutoReject Only	$0.582 \pm 0.143$	$0.700 \pm 0.099$
PyPREP (RANSAC)	$0.557 \pm 0.084$	$0.608 \pm 0.088$
Raw Filter Only	$0.548 \pm 0.101$	$0.599 \pm 0.109$
ICA Only	$0.545 \pm 0.101$	$0.594 \pm 0.115$
ICA + ASR	$0.527 \pm 0.113$	$0.585 \pm 0.147$
ASR Only	$0.524 \pm 0.120$	$0.570 \pm 0.161$
ICA + AutoReject	$0.523 \pm 0.136$	$0.614 \pm 0.157$
ASR + AutoReject	$0.508 \pm 0.149$	$0.624 \pm 0.148$
ICA + ASR + AutoReject	$0.504 \pm 0.123$	$0.583 \pm 0.148$

TABLE IV: Performance of 8- and 16-channel EEG configurations

Model	8 Channels		16 Channels	
	Macro-F1	Bal. Acc	Macro-F1	Bal. Acc
EEGConformer	0.837	0.857	0.901	0.878
EEGNet	0.636	0.662	0.820	0.777
GRU	0.759	0.782	0.870	0.840
ShallowConvNet	0.833	0.835	0.902	0.860

Note: Standard deviations across all entries range from 2.7% to 6.8%.

sensorimotor integration [30], [31]. Across all architectures, increasing the number of channels from 8 to 16 improved accuracy and Macro-F1, indicating that additional spatial information enhances intention-related EEG representations. Although the 8-channel configuration produced competitive results, particularly for convolutional models, the 16-channel setup consistently achieved better generalisation and class discrimination. Consequently, the 16-channel configuration is used for all subsequent experiments.

### D. Evaluation of Model Performance

This section presents a systematic comparison of twelve aforementioned deep learning architectures (Section III-E) for EEG-based driver intention classification. All models were trained using the previously selected configuration.

1) *Overall Performance Across Prediction Horizons:* TABLE V summarises model performance averaged across all prediction horizons (0–1000 ms). The top-performing models (TSception, CNN1D, EEGConformer and ShallowConvNet) achieved the highest Macro-F1 scores, with  $F1 \geq 0.870$ , and demonstrated robust performance across prediction horizons. Among them, TSception attained the best overall results ( $F1 = 0.901 \pm 0.003$ ), followed by CNN1D and EEGConformer. Notably, TSception exhibited exceptional stability across all horizons, indicating strong robustness to temporal offsets from action onset. Several models (DGCNN, GRU, STNet, CCNN, LSTM, and EEGNet) achieved moderate performance ( $0.762 \leq F1 \leq 0.834$ ). While these models provided reliable decoding, they were consistently outperformed by architectures designed for multi-scale temporal-spectral modelling. In contrast, DeepConvNet and ViT exhibited substantially lower Macro-F1 scores. Notably, while ViT achieved a moderate average classification accuracy of (0.588), its

TABLE V: Model performance averaged across all prediction horizons

Ref.	Model	Macro-F1	Bal. Acc
[17]	TSCeption	<b>0.901 ± 0.003</b>	<b>0.907 ± 0.003</b>
[22]	CNN1D	0.894 ± 0.002	0.899 ± 0.002
[27]	EEGConformer	0.874 ± 0.004	0.887 ± 0.004
[18]	ShallowConvNet	0.870 ± 0.005	0.869 ± 0.005
[25]	LSTM	0.834 ± 0.008	0.843 ± 0.012
[25]	GRU	0.830 ± 0.005	0.841 ± 0.004
[21]	DGCNN	0.814 ± 0.004	0.799 ± 0.009
[23]	STNet	0.805 ± 0.003	0.814 ± 0.006
[20]	CCNN	0.803 ± 0.034	0.813 ± 0.045
[19]	EEGNet	0.762 ± 0.004	0.780 ± 0.004
[18]	DeepConvNet	0.641 ± 0.012	0.472 ± 0.017
[28]	ViT	0.572 ± 0.007	0.588 ± 0.019

TABLE VI: Model performance at baseline (0ms) and optimal prediction horizon

Model	Baseline (0ms)		Optimal Horizon		
	Macro-F1	Bal. Acc.	Macro-F1	Bal. Acc.	Time
TSCeption	0.899	0.905	0.905	0.914	400ms
CNN1D	0.894	0.899	0.899	0.905	400ms
EEGConformer	0.870	0.882	0.881	0.891	600ms
ShallowConvNet	0.868	0.868	0.877	0.876	400ms
LSTM	0.835	0.843	0.843	0.857	500ms
GRU	0.835	0.846	0.837	0.848	400ms
DGCNN	0.811	0.795	0.820	0.810	600ms
STNet	0.802	0.808	0.811	0.815	500ms
CCNN	0.755	0.771	0.849	0.856	1000ms
EEGNet	0.761	0.778	0.769	0.786	400ms
DeepConvNet	0.642	0.479	0.666	0.496	900ms
ViT	0.575	0.601	0.585	0.615	800ms

Note: Standard deviations across all entries range from  $\pm 0.03$  to 0.23.

low Macro-F1 score (0.572) indicates poor recognition of minority intention classes, suggesting limitations in handling class imbalance despite reasonable overall accuracy.

2) *Baseline and Optimal Horizon Performance*: To further characterise model behaviour, TABLE VI reports performance at the baseline (0 ms) and at each model’s optimal prediction horizon. TSCeption achieved the strongest baseline performance ( $F1 = 0.899$ ) and peaked at a 400 ms horizon. CNN1D, EEGConformer and ShallowConvNet followed closely, showing robust performance with modest gains at future horizons. EEGConformer also achieved competitive results, with ( $F1 = 0.881$ ) peaking at 600 ms, highlighting the benefit of attention mechanisms for modelling temporally distributed EEG patterns.

Recurrent architectures (GRU and LSTM) and EEGNet achieved moderate performance, with consistent but limited improvements at future horizons, typically peaking around 400-500 ms. In contrast, general-purpose architectures such as ViT and DeepConvNet showed poor generalisation, underscoring the importance of domain-specific inductive biases for EEG-based intention prediction.

### E. Temporal Prediction and Horizon Analysis

Early intention prediction was evaluated across horizons from 0 to 1000 ms ahead of action execution. Fig. 4 shows Macro-F1 trajectories for the top-performing models.

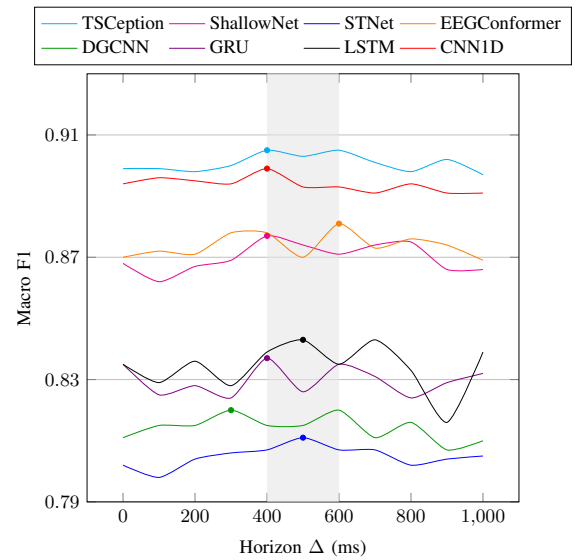


Fig. 4: Macro F1-scores for top 8 models. The grey shaded area (400–600 ms) indicates the neural preparatory window where optimal performance is observed.

Top-performing architectures exhibited strong temporal stability. TSCeption peaked at 400 ms ( $F1 = 0.905$ ) and showed less than 1% degradation at 1000 ms. ShallowConvNet maintained near-baseline performance across all horizons, while EEGConformer showed the largest improvement at future horizons, achieving a +1.26% gain at 600 ms.

Across the top eight models, mean F1 remained stable between 0.852 at 0 ms and 0.857 at 400 ms, with the 400–600 ms range consistently yielding optimal performance. This window aligns with neuroscientific evidence that the late phase of the Bereitschaftspotential (readiness potential) reflects motor-intention formation several hundred milliseconds prior to movement execution [32]. Notably, the majority of top-performing models exhibited less than 1.5% degradation when predicting intentions 1s ahead, demonstrating suitability for anticipatory vehicle control.

### F. Limitations of the Proposed Work

While the proposed framework demonstrates strong performance, several limitations should be noted. Generalisation across unseen subjects, routes, and vehicles was not explicitly evaluated. The dataset is limited to a small number of participants and a predefined route, which may introduce subject- and environment-specific biases. Future work will focus on increasing participants and conducting cross-subject and cross-environment evaluations to improve generalisability.

## V. CONCLUSION

This study established a robust EEG-based framework for driver intention prediction, demonstrating that methodological choices in data preparation and model architecture are paramount. A 1.0-second windowing strategy with label-stratified temporal splitting significantly enhances reliability, while the systematic evaluation of twelve architectures revealed the superiority of domain-specific convolutional mod-

els, with TSception achieving the best performance ( $F1 = 0.901$ ), followed by CNN1D ( $F1 = 0.894$ ) and EEGConformer ( $F1 = 0.874$ ). Minimal preprocessing consistently outperformed aggressive artefact-handling, indicating that raw EEG signals retain essential discriminative features in real-world conditions.

Results demonstrate stable long-horizon prediction, with optimal performance in the 400–600 ms window and less than 1.5% degradation up to 1000 ms. Increasing spatial resolution to 16 channels further improved performance, supporting the feasibility of integrating EEG-based decoding into proactive safety systems.

However, generalisation across unseen subjects, routes, and vehicles remains untested, and the limited number of participants may introduce bias. Future work will address cross-subject evaluation, real-time ADAS integration, and testing under imbalanced data distributions.

## REFERENCES

- [1] Y. Xia, N. Xie, L. Li, P. Angeloudis, D.-H. Lee, and X. M. Chen, "Driving intention understanding for autonomous vehicles: Current approaches and beyond," *Device*, vol. 2, no. 10, 2024.
- [2] X. Liang, Y. Yu, Y. Liu, K. Liu, Y. Liu, and Z. Zhou, "EEG-based emergency braking intention detection during simulated driving," *BioMedical Engineering OnLine*, vol. 22, no. 1, p. 65, 2023.
- [3] P. Liu, T. Qu, H. Gao, and X. Gong, "Driving intention recognition of surrounding vehicles based on a time-sequenced weights hidden markov model for autonomous driving," *Sensors*, vol. 23, no. 21, p. 8761, 2023.
- [4] J. Shi, Y. Lin, Y. Hua, Z. Wang, Z. Zhang, W. Zheng, Y. Song, K. Lu, and S. Lu, "Multi-scenario highway lane-change intention prediction: A physics-informed ai framework for three-class classification," *arXiv preprint arXiv:2509.17354*, 2025.
- [5] O. P. Idowu, A. E. Ilesanmi, X. Li, O. W. Samuel, P. Fang, and G. Li, "An integrated deep learning model for motor intention recognition of multi-class EEG signals in upper limb amputees," *Computer methods and programs in biomedicine*, vol. 206, p. 106121, 2021.
- [6] Y. Ding, C. Udompanyawit, Y. Zhang, and B. He, "EEG-based brain-computer interface enables real-time robotic hand control at individual finger level," *Nature Communications*, vol. 16, no. 1, pp. 1–20, 2025.
- [7] J. Li, Y. Liu, D. Hu, J. Sun, T. Zhou, L. Xu, and Y. Li, "Prediction of driver intention using EEG and GSR signals for advanced driver assistance systems," in *Proceedings of the International Conference on Image Processing, Machine Learning and Pattern Recognition*, 2024, pp. 355–363.
- [8] X. Tao, D. Gao, W. Zhang, T. Liu, B. Du, S. Zhang, and Y. Qin, "A multimodal physiological dataset for driving behaviour analysis," *Scientific data*, vol. 11, no. 1, p. 378, 2024.
- [9] S. Coelli, A. Calcagno, C. M. Cassani, F. Temporiti, P. Reali, R. Gatti, M. Galli, and A. M. Bianchi, "Selecting methods for a modular EEG pre-processing pipeline: An objective comparison," *Biomedical Signal Processing and Control*, vol. 90, p. 105830, 2024.
- [10] N. A. Lutes, V. S. S. Nadendla, and K. Krishnamurthy, "Few-shot transfer learning for individualized braking intent detection on neuromorphic hardware," *Journal of Neural Engineering*, vol. 22, no. 1, p. 016041, 2025.
- [11] Y. Duan, Z. Zhuang, J. Zhou, Y.-C. Chang, Y.-K. Wang, and C.-T. Lin, "Enhancing end-to-end autonomous driving systems through synchronized human behavior data," in *Proceedings of the 1st International Workshop on Brain-Computer Interfaces (BCI) for Multimedia Understanding*, 2024, pp. 1–8.
- [12] OpenBCI, "Gelfree electrode cap guide," <https://docs.openbci.com/AddOns/Headwear/GelfreeElectrodeCap/> (Accessed 28/1/2026), 2026.
- [13] N. Bigdely-Shamlo, T. Mullen, C. Kothe, K.-M. Su, and K. A. Robbins, "The PREP pipeline: standardized preprocessing for large-scale EEG analysis," *Frontiers in neuroinformatics*, vol. 9, p. 16, 2015.
- [14] D. E. Callan, J. J. Torre-Tresols, J. Laguerta, and S. Ishii, "Shredding artifacts: extracting brain activity in EEG from extreme artifacts during skateboarding using ASR and ICA," *Frontiers in neuroergonomics*, vol. 5, p. 1358660, 2024.
- [15] M. Plechawska-Wójcik, P. Augustynowicz, M. Kaczorowska, E. Zabiłska-Mendyk, and D. Zapała, "The influence assessment of artifact subspace reconstruction on the EEG signal characteristics," *Applied Sciences*, vol. 13, no. 3, p. 1605, 2023.
- [16] S. Iwama, S. Tsuchimoto, N. Mizuguchi, and J. Ushiba, "EEG decoding with spatiotemporal convolutional neural network for visualization and closed-loop control of sensorimotor activities: A simultaneous EEG-fMRI study," *Human Brain Mapping*, vol. 45, no. 9, p. e26767, 2024.
- [17] Y. Ding, N. Robinson, S. Zhang, Q. Zeng, and C. Guan, "TSception: Capturing temporal dynamics and spatial asymmetry from EEG for emotion recognition," *IEEE Transactions on Affective Computing*, vol. 14, no. 3, pp. 2238–2250, 2022.
- [18] R. T. Schirmeister, J. T. Springenberg, L. D. J. Fiederer, M. Glasstetter, K. Eggensperger, M. Tangermann, F. Hutter, W. Burgard, and T. Ball, "Deep learning with convolutional neural networks for EEG decoding and visualization," *Human brain mapping*, vol. 38, no. 11, pp. 5391–5420, 2017.
- [19] V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, "EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces," *Journal of neural engineering*, vol. 15, no. 5, p. 056013, 2018.
- [20] Y. Yang, Q. Wu, Y. Fu, and X. Chen, "Continuous convolutional neural network with 3D input for EEG-based emotion recognition," in *International conference on neural information processing*. Springer, 2018, pp. 433–443.
- [21] T. Song, W. Zheng, P. Song, and Z. Cui, "EEG emotion recognition using dynamical graph convolutional neural networks," *IEEE Transactions on Affective Computing*, vol. 11, no. 3, pp. 532–541, 2018.
- [22] M. Taghizadeh, F. Vaez, and M. Faezipour, "EEG motor imagery classification by feature extracted deep 1D-CNN and semi-deep fine-tuning," *IEEE Access*, 2024.
- [23] Z. Zhang, Y. Liu, and S.-h. Zhong, "GANSER: A self-supervised data augmentation framework for EEG-based emotion recognition," *IEEE Transactions on Affective Computing*, vol. 14, no. 3, pp. 2048–2063, 2022.
- [24] A. Sikka, H. Jamalabadi, M. Krylova, S. Alizadeh, J. N. Van Der Meer, L. Danyeli, M. Deliano, P. Vicheva, T. Hahn, T. Koenig *et al.*, "Investigating the temporal dynamics of electroencephalogram EEG microstates using recurrent neural networks," *Human brain mapping*, vol. 41, no. 9, pp. 2334–2346, 2020.
- [25] X. Zhang and L. Yao, *Deep learning for EEG-based brain-computer interfaces: representations, algorithms and applications*. World Scientific, 2021.
- [26] X. Ma, W. Chen, Z. Pei, Y. Zhang, and J. Chen, "Attention-based convolutional neural network with multi-modal temporal information fusion for motor imagery EEG decoding," *Computers in Biology and Medicine*, vol. 175, p. 108504, 2024.
- [27] Y. Song, Q. Zheng, B. Liu, and X. Gao, "EEG Conformer: Convolutional transformer for EEG decoding and visualization," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 31, pp. 710–719, 2022.
- [28] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [29] J. Yoon, D. Jarrett, and M. Van der Schaar, "Time-series generative adversarial networks," *Advances in neural information processing systems*, vol. 32, 2019.
- [30] Z. Khalilidali, R. Chavarriaga, L. A. Gheorghe, and J. del R Millán, "Action prediction based on anticipatory brain potentials during simulated driving," *Journal of neural engineering*, vol. 12, no. 6, p. 066006, 2015.
- [31] S. Haufe, M. S. Treder, M. F. Gugler, M. Sagebaum, G. Curio, and B. Blankertz, "EEG potentials predict upcoming emergency brakings during simulated driving," *Journal of neural engineering*, vol. 8, no. 5, p. 056001, 2011.
- [32] R. Verleger, M. Haake, A. Baur, and K. Śmigasiwicz, "Time to move again: does the Bereitschaftspotential covary with demands on internal timing?" *Frontiers in human neuroscience*, vol. 10, p. 642, 2016.